



page size dependency

Found 7,430 of 243,132

Terms used: [page size dependency](#)

Sort results by relevance

Display results in expanded form

Save results to a Binder

Open results in a new window

Refine these results with [Advanced Search](#)Try this search in [The ACM Guide](#)

Results 1 - 20 of 7,430

Result page: 1 2 3 4 5 6 7 8 9 10 [next](#) [>>](#)

1 Reducing coherence overhead of barrier synchronization in software DSMs

Jae Bum Lee, Chu Shik Jhon

November 1998 Supercomputing '98: Proceedings of the 1998 ACM/IEEE conference on Supercomputing (CDROM)

Publisher: IEEE Computer Society

Full text available: [pdf\(306.05 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Bibliometrics: Downloads (6 Weeks): 1, Downloads (12 Months): 9, Citation Count: 0

Software Distributed Shared Memory (SDSM) systems usually have the large coherence granularity that is imposed by the underlying virtual memory page size. To alleviate the coherence overheads such as the net worktraffic to preserve the coherence, or page ...

Keywords: automatic-update, data coherence, data dependency analysis, relaxed memory models, run-time support, software distributed shared memory

2 Properties of the working-set model

Peter J. Denning, Stuart C. Schwartz

June 1972 ACM SIGOPS Operating Systems Review, Volume 6 Issue 1/2

Publisher: ACM

Full text available: [pdf\(36.04 KB\)](#) Additional Information: [full citation](#), [abstract](#)

Bibliometrics: Downloads (6 Weeks): 0, Downloads (12 Months): 12, Citation Count: 0

A program's working set $W(t, T)$ at time t is the set of distinct pages among the T most recently referenced pages. Relations between the average working-set size, the missing-page rate, and the interreference-interval distribution may be derived both ...

Ads by Google

Find Your Match - Free

Browse Photos and Meet Singles Free
Search and Profiles[www.True.com](#)

Real-time Data Processing Monitor, Analyze, Act on real-time data streams as they happen.

[www.Progress.com/Apam](#)

550 Credit Score

The Average US Credit Score is 675. See Your Score for \$0.

[www.CreditScore.name](#)

Avg. US Credit Score: 678

The Cost to see yours is: \$0 Easy to Read and Viewable Online

[www.FreeCreditReport.com](#)

3 An instant and accurate size estimation method for joins and selections in a retrieval-intensive environment

Wei Sun, Yibei Ling, Naphtali Rische, Yi Deng

June 1993 ACM SIGMOD Record, Volume 22 Issue 2

Publisher: ACM

Full text available: [pdf\(1.05 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [Index terms](#)

Bibliometrics: Downloads (6 Weeks): 3, Downloads (12 Months): 29, Citation Count: 19

This paper proposes a novel strategy for estimating the size of the resulting relation after an equi-join and selection using a regression model. An approximating series representing the underlying data distribution and dependency is derived ...

4 Properties of the working-set model

◆ Peter J. Denning, Stuart C. Schwartz

◆ March 1972 Communications of the ACM, Volume 15 Issue 3

Publisher: ACM

Full text available:  pdf(672.49 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#)

Bibliometrics: Downloads (6 Weeks): 8, Downloads (12 Months): 79, Citation Count: 45

A program's working set $W(t, T)$ at time t is the set of distinct pages among the T most recently referenced pages. Relations between the average working-set size, the missing-page rate, ...

Keywords: paging, paging algorithms, program behavior, program modeling, working-set model

5 An instant and accurate size estimation method for joins and selections in a

◆ retrieval-intensive environment

◆ Wei Sun, Yibei Ling, Naphtali Rishe, Yi Deng

June 1993 SIGMOD '93: Proceedings of the 1993 ACM SIGMOD international conference on Management of data

Publisher: ACM

Full text available:  pdf(1.05 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 3, Downloads (12 Months): 29, Citation Count: 19

This paper proposes a novel strategy for estimating the size of the resulting relation after an equi-join and selection using a regression model. An approximating series representing the underlying data distribution and dependency is derived ...

6 Properties of the working-set model

◆ Peter J. Denning, Stuart C. Schwartz

◆ October 1971 SOSP '71: Proceedings of the third ACM symposium on Operating systems principles

Publisher: ACM

Full text available:  pdf(28.37 KB) Additional Information: [full citation](#), [abstract](#), [cited by](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 1, Downloads (12 Months): 5, Citation Count: 1

This paper appears in the March, 1972, issue of the Communications of the ACM. Its abstract is reproduced below. A program's working set $W(t, T)$ at time t is the set of distinct pages among the T most recently referenced pages. Relations between the average ...

7 Proceedings of the 17th international conference on Computational linguistics -

Volume 2

Christian Boitet, Pete Whitelock

August 1998 proceeding Volume 2

Publisher: Association for Computational Linguistics

Additional Information: [full citation](#), [abstract](#)

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citation Count: 0

This volume may come as a surprise to regular participants of the annual ACL conferences. It is certainly more substantial in its weight and, hopefully, in its contents as well. While the former is chiefly the result of the extremely difficult and ...

8 Proceedings of the 36th annual meeting on Association for Computational

Linguistics - Volume 1

Christian Boitet, Pete Whitelock

August 1998 proceeding Volume 1

Publisher: Association for Computational Linguistics

Additional Information: [full citation](#), [abstract](#)

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citation Count: 0

This volume may come as a surprise to regular participants of the annual ACL conferences. It is certainly more substantial in its weight and, hopefully, in its contents as well. While the former is chiefly the result of the extremely difficult and ...

9 Proceedings of the 17th international conference on Computational linguistics - Volume 1

Christian Boitet, Pete Whitelock

August 1998 proceeding Volume 1

Publisher: Association for Computational Linguistics

Additional Information: [full citation](#), [abstract](#)

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citation Count: 0

This volume may come as a surprise to regular participants of the annual ACL conferences. It is certainly more substantial in its weight and, hopefully, in its contents as well. While the former is chiefly the result of the extremely difficult and ...

10 Proceedings of the 36th annual meeting on Association for Computational Linguistics - Volume 2

Christian Boitet, Pete Whitelock

August 1998 proceeding Volume 2

Publisher: Association for Computational Linguistics

Additional Information: [full citation](#), [abstract](#)

Bibliometrics: Downloads (6 Weeks): n/a, Downloads (12 Months): n/a, Citation Count: 0

This volume may come as a surprise to regular participants of the annual ACL conferences. It is certainly more substantial in its weight and, hopefully, in its contents as well. While the former is chiefly the result of the extremely difficult and ...

11 Compiler-based I/O prefetching for out-of-core applications

 Angela Demke Brown, Todd C. Mowry, Orran Krieger

May 2001 ACM Transactions on Computer Systems (TOCS), Volume 19 Issue 2

Publisher: ACM

Full text available:  [pdf\(499.93 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#), [review](#)

Bibliometrics: Downloads (6 Weeks): 7, Downloads (12 Months): 79, Citation Count: 11

Current operating systems offer poor performance when a numeric application's working set does not fit in main memory. As a result, programmers who wish to solve "out-of-core" problems efficiently are typically faced with the onerous task ...

Keywords: compiler optimization, prefetching, virtual memory

12 Proxy-based acceleration of dynamically generated content on the world wide web:

 **An approach and implementation**

Anindya Datta, Kaushik Dutta, Helen Thomas, Debra Vandermeer, Krithi Ramamritham

June 2004 ACM Transactions on Database Systems (TODS), Volume 29 Issue 2

Publisher: ACM

Full text available:  pdf(927.23 KB) Additional Information: [full citation](#), [appendices and supplements](#), [abstract](#), [references](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 216, Citation Count: 1

As Internet traffic continues to grow and websites become increasingly complex, performance and scalability are major issues for websites. Websites are increasingly relying on dynamic content generation applications to provide website visitors with dynamic, ...

Keywords: Edge caching, caching dynamically generated content, fragment caching, implementation, proxy caching, world wide web

13 GPGPU: general purpose computation on graphics hardware

 David Luebke, Mark Harris, Jens Krüger, Tim Purcell, Naga Govindaraju, Ian Buck, Cliff Woolley, Aaron Lefohn

August 2004 SIGGRAPH '04: ACM SIGGRAPH 2004 Course Notes

Publisher: ACM

Full text available:  pdf(63.03 MB) Additional Information: [full citation](#), [abstract](#), [cited by](#)

Bibliometrics: Downloads (6 Weeks): 105, Downloads (12 Months): 1438, Citation Count: 3

The graphics processor (GPU) on today's commodity video cards has evolved into an extremely powerful and flexible processor. The latest graphics architectures provide tremendous memory bandwidth and computational horsepower, with fully programmable vertex ...

14 I/O reference behavior of production database workloads and the TPC benchmarks—an analysis at the logical level

 Windsor W. Hsu, Alan Jay Smith, Honesty C. Young

March 2001 ACM Transactions on Database Systems (TODS), Volume 26 Issue 1

Publisher: ACM

Full text available:  pdf(5.42 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 11, Downloads (12 Months): 156, Citation Count: 8

As improvements in processor performance continue to far outpace improvements in storage performance, I/O is increasingly the bottleneck in computer systems, especially in large database systems that manage huge amounts of data. The key to achieving ...

Keywords: I/O, TPC benchmarks, caching, locality, prefetching, production database workloads, reference behavior, sequentiality, workload characterization

15 Path and cache conscious prefetching (PCOP)

Zhen He, Alonso Marquez

April 2007 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 16 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(582.94 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 92, Citation Count: 0

Main memory cache performance continues to play an important role in determining the overall performance of object-oriented, object-relational and XML databases. An effective method of improving main memory cache performance is to prefetch or pre-load ...

Keywords: Caching, Clustering, Databases, Prefetching

◆ Effect of node size on the performance of cache-conscious B⁺-trees

Richard A. Hankins, Jignesh M. Patel

June 2003 ACM SIGMETRICS Performance Evaluation Review, Volume 31 Issue 1

Publisher: ACM

Full text available:  pdf(271.16 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 70, Citation Count: 1

In main-memory databases, the number of processor cache misses has a critical impact on the performance of the system. Cache-conscious indices are designed to improve performance by reducing the number of processor cache misses that are incurred during ...

Keywords: B⁺-tree, cache-conscious, index**17 Effect of node size on the performance of cache-conscious B⁺-trees**

Richard A. Hankins, Jignesh M. Patel

June 2003 SIGMETRICS '03: Proceedings of the 2003 ACM SIGMETRICS international conference on Measurement and modeling of computer systems

Publisher: ACM

Full text available:  pdf(271.16 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 70, Citation Count: 1

In main-memory databases, the number of processor cache misses has a critical impact on the performance of the system. Cache-conscious indices are designed to improve performance by reducing the number of processor cache misses that are incurred during ...

Keywords: B⁺-tree, cache-conscious, index**18 Data page layouts for relational databases on deep memory hierarchies**

Anastassia Ailamaki, David J. DeWitt, Mark D. Hill

November 2002 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 11 Issue 3

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(593.86 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 74, Citation Count: 3

Relational database systems have traditionally optimized for I/O performance and organized records sequentially on disk pages using the N-ary Storage Model (NSM) (a.k.a., slotted pages). Recent research, however, indicates that cache utilization ...

Keywords: Cache-conscious database systems, Disk page layout, Relational data placement

19 Dynamic vp-tree indexing for *n*-nearest neighbor search given pair-wise distances

Ada Wai-chee Fu, Polly Mei-shuen Chan, Yin-Ling Cheung, Yiu Sang Moon

July 2000 The VLDB Journal — The International Journal on Very Large Data Bases, Volume 9 Issue 2

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(232.09 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#)

Bibliometrics: Downloads (6 Weeks): 5, Downloads (12 Months): 66, Citation Count: 10

For some multimedia applications, it has been found that domain objects cannot be represented as feature vectors in a multidimensional space. Instead, pair-wise distances between data objects are the only input. To support content-based retrieval, one ...

Keywords: Content-based retrieval, Indexing, Nearest neighbor search, Pair-wise distances, Updating

20 A cost model for query processing in high dimensional data spaces

Christian Böhm

June 2000 ACM Transactions on Database Systems (TODS), Volume 25 Issue 2
Publisher: ACM

Full text available: [pdf \(362.22 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [cited by](#), [index terms](#), [review](#)

Bibliometrics: Downloads (6 Weeks): 6, Downloads (12 Months): 108, Citation Count: 17

During the last decade, multimedia databases have become increasingly important in many application areas such as medicine, CAD, geography, and molecular biology. An important research topic in multimedia databases is similarity search in large data ...

Keywords: cost model, multidimensional index

Results 1 - 20 of 7,430

Result page: 1 2 3 4 5 6 7 8 9 10 next >>

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)